

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 271 953 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
02.01.2003 Bulletin 2003/01

(51) Int Cl. 7: H04N 7/24

(21) Application number: 02011057.3

(22) Date of filing: 17.05.2002

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR

Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 28.06.2001 US 895872

(71) Applicant: MICROSOFT CORPORATION
Redmond, Washington 98052-6399 (US)

(72) Inventors:

• Colville, Scott, E.
Kirkland, Washington 98033 (US)

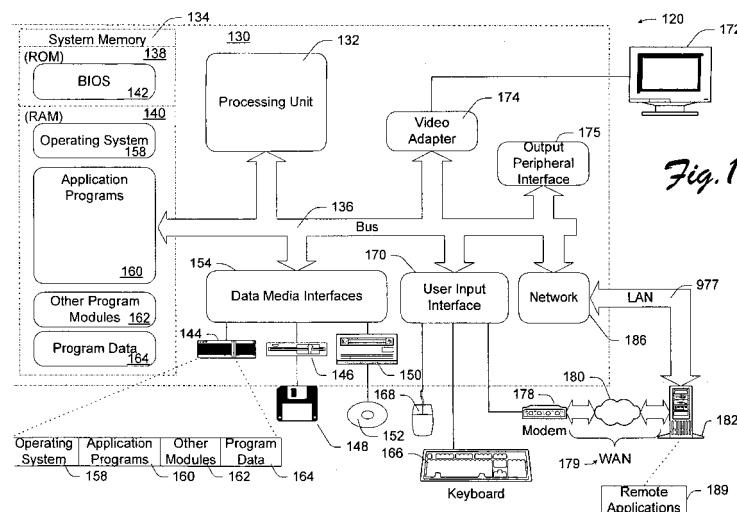
• Bhatt, Sanjay
Redmond, Washington 98052 (US)
• Klemets, Anders E.
Seattle, Washington 98122 (US)
• Batterberry, Troy
Kirkland, Washington 98033 (US)

(74) Representative: Grünecker, Kinkeldey,
Stockmair & Schwanhäusser Anwaltssozietät
Maximilianstrasse 58
80538 München (DE)

(54) Improved startup methods and apparatuses for use in streaming content

(57) Methods and apparatuses are provided for use with a client and server device connected through a communication link. The client device sends a startup request to the server device. The startup request identifies a streamable media content that is to be provided to the client device, a communication link bandwidth associated with the communication link, and an amount of the desired streamable media content that is to be provided at a bitrate greater than the encoded bitrate but no greater than about the communication link bandwidth. The server device buffers at least the amount of

the streamable media content, and transmits the amount of the buffered streamable media content at the higher bitrate. The server device locates a discrete rendering point in the amount of the buffered streamable media content and initiates transmission beginning with the discrete rendering point. After transmitting the amount of the buffered streamable media content, the server device transmits subsequent portions of the streamable media content to the client device at a bitrate about equal to the encoded bitrate. The client device buffers received streamable media content, and subsequently renders the buffered streamed media content.



Description

RELATED APPLICATIONS

[0001] The present application is related to co-pending U.S. Patent Application No. 09/636,004, filed August 9, 2000, and titled "Fast Dynamic Measurement of Connection Bandwidth", which is incorporated by reference herein.

TECHNICAL FIELD

[0002] The present invention relates generally to streaming media devices, and more particularly to methods and apparatuses that provide improved end-user startup times when streaming content.

BACKGROUND

[0003] Today, nearly every computer user is well accustomed to the broadcast television medium. When a new television channel is requested, the new channel is generally displayed (rendered) very quickly. The same can be said for conventional broadcast radio stations.

[0004] Unfortunately, the same cannot be said for conventional Internet streamed content. Streaming media is typically provided from a server device to a client device over the Internet or other like network. For a variety of technical reasons, the end-user experience can be degraded, for example, by pauses experienced in the rendering due to late-arriving/recent data. Such pauses, however, tend to occur randomly and in certain configurations occur very rarely. However, there is one pause that most end-users experience every time they select a streaming media program, namely, a slow startup time.

[0005] This poor startup experience tends to inhibit the adoption of streaming media in many markets. It is also tends to reduce the amount of time end-users are willing to use the technology. Thus, channel "surfing" is largely unacceptable with conventional streaming techniques. Hence, there is a need for improved streaming media methods and apparatuses that can significantly reduce the startup time that the end-user experiences.

SUMMARY

[0006] In accordance with certain aspects of the present invention, improved streaming media methods and apparatuses are provided that significantly reduce the startup time that the end-user experiences.

[0007] By way of example, the above stated needs and others are met by a system in accordance with certain implementations of the present invention. The system includes a client device and a server device, which are operatively connected through a communication link. The client device is configured to send at least one startup request to the server device over the communica-

cation link. The startup request identifies a streamable media content that is to be provided to the client device, a communication link bandwidth associated with the communication link, and an amount of the desired streamable media content that is to be provided at a bitrate greater than the encoded bitrate, but no greater than about the communication link bandwidth. The server device is configured to buffer at least the amount of the streamable media content and transmit the amount of the buffered streamable media content at the higher bitrate. After transmitting the amount of the buffered streamable media content, the server device transmits subsequent portions of the streamable media content to the client device at a bitrate about equal to the encoded bitrate. The client device is configured to buffer received streamable media content, and subsequently render the buffered streamed media content.

[0008] In accordance with certain implementations, the server device locates a discrete rendering point in the amount of the buffered streamable media content and initiates transmission beginning with the discrete rendering point.

[0009] In accordance with certain further implementations, the client device determines the communication link bandwidth.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] A more complete understanding of the various methods and apparatuses of the present invention may be had by reference to the following detailed description when taken in conjunction with the accompanying drawings wherein:

35 Fig. 1 is a block diagram depicting an exemplary device, in the form of a computer, which is suitable for use in providing, receiving, and/or otherwise communicating streamed media, in accordance with certain implementations of the present invention.

40 Fig. 2 is a block diagram depicting an exemplary streaming media system having a server device and a client device, in accordance with certain implementations of the present invention.

45 Fig. 3 is an illustrative diagram depicting an exemplary content stream, suitable for streaming in the streaming media system of Fig. 2, in accordance with certain implementations of the present invention.

50 Fig. 4 is a flow diagram depicting an exemplary client-centric media streaming process suitable for use in the client device of Fig. 2, in accordance with certain implementations of the present invention.

55 Fig. 5 is a flow diagram depicting an exemplary fast startup media streaming process suitable for use in the server device of Fig. 2, in accordance with certain implementations of the present invention.

Fig. 6 is a time line diagram depicting the transmis-

sion data bitrate for an exemplary fast startup streaming media transmission associated with the streaming media system of Fig. 2, in accordance with certain implementations of the present invention.

DETAILED DESCRIPTION

[0011] Turning to the drawings, wherein like reference numerals refer to like elements, the invention is illustrated as being implemented in a suitable computing environment. Although not required, portions of the invention are described in the general context of computer-executable instructions, such as program modules, being executed by a computer or like device, which, for example, may take the form of a personal computer (PC), a workstation, a portable computer, a server, a plurality of processors, a mainframe computer, a wireless communications base station, a hand-held communications device, a streamed media player, a set-top box, etc.

[0012] Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The various exemplary implementations of the present invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0013] As provided herein, the term "logic" is meant to apply to any form of logic and requisite supporting elements, including, e.g., software, firmware, hardware, and/or any combination thereof.

[0014] Fig. 1 illustrates an example of a suitable computing environment 120 on which portions of the subsequently described methods and apparatuses may be implemented.

[0015] Exemplary computing environment 120 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the improved methods and apparatuses described herein. Neither should computing environment 120 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in computing environment 120.

[0016] The improved methods and apparatuses herein are operational with numerous other general purpose and/or special purpose computing system environments or configurations.

[0017] As shown in Fig. 1, computing environment 120 includes a general-purpose computing device in the form of a computer 130. The components of computer 130 may include one or more processors or processing units 132, a system memory 134, and a bus 136 that couples various system components including system

memory 134 to processor 132.

[0018] Bus 136 represents one or more of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnects (PCI) bus also known as Mezzanine bus.

[0019] Computer 130 typically includes a variety of computer readable media. Such media may be any available media that is accessible by computer 130, and it includes both volatile and non-volatile media, removable and non-removable media.

[0020] In Fig. 1, system memory 134 includes computer readable media in the form of volatile memory, such as random access memory (RAM) 140, and/or non-volatile memory, such as read only memory (ROM) 138. A basic input/output system (BIOS) 142, containing the basic routines that help to transfer information between elements within computer 130, such as during start-up, is stored in ROM 138. RAM 140 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processor 132.

[0021] Computer 130 may further include other removable/non-removable, volatile/non-volatile computer storage media. For example, Fig. 1 illustrates a hard disk drive 144 for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"), a magnetic disk drive 146 for reading from and writing to a removable, non-volatile magnetic disk 148 (e.g., a "floppy disk"), and an optical disk drive 150 for reading from or writing to a removable, non-volatile optical disk 152 such as a CD-ROM, CD-R, CD-RW, DVD-ROM, DVD-RAM or other optical media.

40 Hard disk drive 144, magnetic disk drive 146 and optical disk drive 150 are each connected to bus 136 by one or more interfaces 154.

[0022] The drives and associated computer-readable media provide nonvolatile storage of computer readable instructions, data structures, program modules, and other data for computer 130. Although the exemplary environment described herein employs a hard disk, a removable magnetic disk 148 and a removable optical disk 152, it should be appreciated by those skilled in the art that other types of computer readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, random access memories (RAMs), read only memories (ROM), and the like, may also be used in the exemplary operating environment.

[0023] A number of program modules may be stored on the hard disk, magnetic disk 148, optical disk 152, ROM 138, or RAM 140, including, e.g., an operating sys-

tem 158, one or more application programs 160, other program modules 162, and program data 164.

[0024] The improved methods and apparatuses described herein may be implemented within operating system 158, one or more application programs 160, other program modules 162, and/or program data 164.

[0025] A user may provide commands and information into computer 130 through input devices such as keyboard 166 and pointing device 168 (such as a "mouse"). Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, serial port, scanner, camera, etc. These and other input devices are connected to the processing unit 132 through a user input interface 170 that is coupled to bus 136, but may be connected by other interface and bus structures, such as a parallel port, game port, or a universal serial bus (USB).

[0026] A monitor 172 or other type of display device is also connected to bus 136 via an interface, such as a video adapter 174. In addition to monitor 172, personal computers typically include other peripheral output devices (not shown), such as speakers and printers, which may be connected through output peripheral interface 175.

[0027] Computer 130 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 182. Remote computer 182 may include many or all of the elements and features described herein relative to computer 130.

[0028] Logical connections shown in Fig. 1 are a local area network (LAN) 177 and a general wide area network (WAN) 179. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet.

[0029] When used in a LAN networking environment, computer 130 is connected to LAN 177 via network interface or adapter 186. When used in a WAN networking environment, the computer typically includes a modem 178 or other means for establishing communications over WAN 179. Modem 178, which may be internal or external, may be connected to system bus 136 via the user input interface 170 or other appropriate mechanism.

[0030] Depicted in Fig. 1, is a specific implementation of a WAN via the Internet. Here, computer 130 employs modem 178 to establish communications with at least one remote computer 182 via the Internet 180.

[0031] In a networked environment, program modules depicted relative to computer 130, or portions thereof, may be stored in a remote memory storage device. Thus, e.g., as depicted in Fig. 1, remote application programs 189 may reside on a memory device of remote computer 182. It will be appreciated that the network connections shown and described are exemplary and other means of establishing a communications link between the computers may be used.

[0032] Reference is now made to Fig. 2, which depicts

an exemplary streaming media system 200 having a server device 202 operatively coupled to a network 204 and configured to stream media there over to a client device 206 that is also operatively coupled to network 204.

[0033] Streaming media system 200 is arranged in a client-centric configuration in which client device 204 selects a streaming media program on server device 202, specifies that a particular fast startup is desired, and provides information to server device 202 about the communication link over which the streaming media will be carried. In this manner, client device 204 is able to control the streaming media process and server device 202. As described in greater detail below, in controlling the streaming media the client device 204 causes server device 202 to stream media during an initial period of time at data bitrate that is greater than the media's encoded bitrate. This allows client device 204 to quickly receive data and begin the rendering process sooner.

[0034] This is unlike previous server-centric solutions used to provide video on-demand, such as, for example, the system and method presented in U.S. Patent No. 5,963,202, issued to Nathaniel Polish. In such server-centric systems, the server device, rather than the client device, has control over a video data transfer. Thus, for example, a server needs to determine how much video data can be transferred over the communication link and when to transfer it during the progressive download. One of the drawbacks to a server-centric system is that the server is required to monitor, for every client device, the status of the communications link and data buffers in the client device. While a progressive video download technique may be efficient for an in-home or hotel video-on-demand system, it would likely prove inefficient in a larger network environment, such as, for example, the Internet, a corporate intranet, a wide area network, a wireless network, etc.

[0035] Returning to Fig. 2, as depicted, server 202 includes a media server 208. Media server 208 includes fast startup logic 210 and is operatively coupled to a buffer 212. As shown, in this example, media server is operatively coupled to a first interface 214 that provides access to a media storage device 216. Media server 208 is further operatively coupled to a second interface 218 that provides access to a broadcast media device 220 (represented by a video camera). Media server 208 is operatively coupled to network 204 through a third interface 222. It is recognized that in other implementations interfaces 214, 218 and/or 222 may be combined in some manner.

[0036] As its name suggests, media server 208 is configured to serve or otherwise provide streaming media to client device 206. To accomplish this task, media server 208 exchanges information with client device 206 through interface 222 and network 204. The techniques and protocols used to provide communications between server device 202 and client device 206 are well known and therefore will not be described in to great of detail.

[0037] Media server 208 identifies the availability of streaming media programs to client device 206. In this example, media server 208 accesses/receives streaming media programs from two sources, namely media storage device 216 and broadcast media device 220. Media storage device 216 is representative of a data storage device, such as, for example, one or more magnetic disk drives, one or more optical disc drives, and the like. Here, media storage device 216 is configured to allow media server 208 to stream media "on-demand" to client device 206. As used herein, "on-demand" means that the media is stored in media storage device 216, and has since then been made available for streaming and replay at subsequent times. Thus, for example, an earlier news program may be recorded and stored in its entirety on media storage device 216 and subsequently made available on-demand.

[0038] To the contrary, broadcast media device 220 is representative of media that has not been significantly stored, and certainly not in its entirety. An example would be a news program that is being aired in real-time or near real-time. Such a media program would not, therefore, be available on-demand. When client device 206 selects this broadcast media stream, the streaming media will "jump" into the news program at about the point where it is being aired.

[0039] Buffer 212 is used by media server 208 to temporarily store media data in support of the streaming process. Buffer 208 would typically include random access memory.

[0040] As shown, client device 206 includes a media client 224. Media client 224 is configured to support the selection, receipt and rendering of streaming media from server device 202 via interface 232 and network 204. To further accomplish its tasks, media client 224 includes fast startup logic 226 and renderer 228. Media client 224 is also operatively coupled to a buffer 230. Buffer 230 typically includes random access memory. Renderer 228 is configured to process the streamed media data and render the data as applicable for client device 206 and the received media. Rendering processes are well known, and the details of such are beyond the scope of the present invention.

[0041] With this exemplary streaming media system in mind, this detailed description will now focus on the functionality of media client 224 and fast startup logic 210 in server device 202 and corresponding fast start logic 226 in client device 206.

[0042] Media client 224 requires buffering of the streaming data for a variety of reasons. For example, buffering allows client device 206 to request and successfully obtain retransmissions when content packets are lost without impacting continuous playback. Buffering also allows playback to be smooth on networks that have jitter or inconsistent bandwidth response. Highly efficient compression technology often requires a significant duration of content (e.g., an entire frame or more) to be present on the client before decompression can

begin. All of these issues contribute to the necessity of buffering content by media client 224.

[0043] Depending on the compression technology and the content type, buffering can vary anywhere from 5 less than 1 second to many seconds. Certain conventional media players, for example, buffer roughly five seconds worth of content before allowing any rendering to begin. Because conventional streaming media servers are designed to deliver the content at the encoded 10 bitrate, the end-user will have to wait at least five seconds for the buffers to fill and rendering to begin.

[0044] Since broadcast media is typically already running when most client devices connect, a client device 15 may be required to wait before even beginning the buffering process. With typical compression technologies in use today, for example, buffering needs to start at certain discrete points in the stream. When an individual client device subscribes to a conventional broadcast stream, it will need to wait for one of these discrete 20 points to appear in the stream before even starting the buffering process. The frequency of the discrete points can vary dramatically depending on the compression technology used, the content type, and even the content characteristics. These discrete buffering points can vary 25 in frequency from several times a second to as little as once every sixteen seconds or less.

[0045] Given that a conventional streaming media client must first wait to find a discrete entry point and then 30 wait for the buffers to fill, the user often experiences significant delay when attempting to start rendering a streaming media broadcast. In accordance with certain aspects of the present invention, methods and apparatuses are provided that tend to significantly reduce the time required to fill the client device's buffer(s) and ultimately 35 allow rendering to begin faster for both broadcast and on-demand content. Thus, for example, in certain exemplary implementations, additional available network bandwidth is utilized to accelerate the streaming of content and as such fill the client device's buffer(s) 40 faster. Hence, the term fast startup.

[0046] In accordance with certain aspects of the 45 present invention, the various fast startup methods and apparatuses can be implemented by extending the usage/syntax of conventional streaming protocols, such as, for example, Microsoft Media Server (MMS), Real Time Streaming Protocol (RTSP), HyperText Transfer Protocol (HTTP), and the like.

[0047] Co-pending U.S. Patent Application No. 09/636,004, filed August 9, 2000, and titled "Fast Dynamic Measurement of Connection Bandwidth", which 50 is incorporated by reference herein, describes, in greater detail, techniques by which media client 224 can determine the bandwidth present between client device 206 and server device 202, prior to requesting the actual 55 delivery of the streaming media. This bandwidth is known as the link bandwidth.

[0048] Basically, the fast dynamic measurement of connection bandwidth utilizes a single pair of packets to

calculate bandwidth between client device 206 and server device 202. This calculation is based upon a packet-pair technique. This bandwidth measurement is extremely quick. On its journey across network 204, communication equipment and modems may compress a packet. This compression shrinks the size of the packet; thus, it can distort the bandwidth calculation using such a shrunken packet. To avoid this distortion, the fast dynamic measurement of connection bandwidth employs non-compressible packets. More specifically, it employs highly entropic packets. Therefore, a packet cannot be compressed during its journey. In addition, on its journey across network 204, packets may be rerouted, delayed, misrouted, and the like. These momentary delays may result in a momentary bad bandwidth calculation. This problem is ameliorated by using a history list (not shown) at media client 224 that keeps track of recent measurements. Media client 224 can then determine the median value from the history list. That median value is representative of the link bandwidth.

[0049] This represents one exemplary technique for determining the link bandwidth. Those skilled in the art will recognize that other techniques may be employed to determine to some degree of certainty the link bandwidth.

[0050] Media client 224 can use conventional protocol, such as, e.g., a session description protocol (SDP) to communicate with media server 208 and identify the location and characteristics of the available streaming media.

[0051] In this manner, media client 224 is therefore able to determine both the link bandwidth and also the bandwidth of the individual stream(s) in the streaming media program. As such, fast startup logic 226 in media client 224 can request that the content be initially streamed at a rate faster than the encoded bitrate of the content. This request for fast startup is handled by fast startup logic 210 in media server 208.

[0052] Assuming normal playback speed, streaming the content at a rate greater than the encoded bitrate implies that the amount of data in client buffer 230 will increase in size over time. It is undesirable to continue to stream the content at a rate faster than the encoded bitrate of the content indefinitely, given the limited amount of memory in buffer 230. Instead, client buffer 230 is sufficiently filled at the fast rate at the beginning of the streaming process, and subsequently the streaming rate changes to roughly match the encoded bitrate of the media program (file). This design has the benefit of using the additional link bandwidth to quickly fill client buffer 230 without requiring additional memory in buffer 230.

[0053] Fast startup logic 210, within media server 208, is configured to respond to the fast startup request by streaming the content at the faster rate. In the case of broadcast media, such as a live video feed, fast startup logic 210 temporarily stores a portion of the streaming broadcast media to server buffer 212. In this manner,

new client devices connecting to server device 202 can be sent content packets at a rate greater than the encoded bitrate of the broadcast stream.

[0054] Thus, for example, in certain implementations 5 if the content is encoded at 16 kbps, then fast startup logic 210 will store the previous 10 seconds of the broadcast media in buffer 212. This exemplary buffering process therefore requires 20 Kbytes of memory. As a result, client devices that connect after the broadcast has started 10 are able to request approximately up to about 10 seconds of content at a rate much faster than 16 kbps.

[0055] This is just one example; in other implementations, the buffering process may store a longer or shorter amount of the broadcast media in buffer 212.

[0056] Startup logic 210 is further configured to intelligently decide where to start sending content packets from buffer 212 as new clients connect to the broadcast. For example, assume that a broadcast program is running and a new client connects to server 202. If startup 15 logic 210 has buffered the previous 10 seconds of content in buffer 212, then theoretically fast startup logic 210 can start sending content at roughly any point from time ConnectTime_{clientX}-10 to ConnectTime_{clientX}.

[0057] However, starting the streaming of content at 20 the beginning of the 10 second buffer can be problematic because the content residing at ConnectTime_{clientX}-10 may not contain a discrete starting point as required by media client 224. Typically, for certain types of 25 streamed content, media client 224 can only start rendering the content at discrete points within the streamed data, such as, for example, certain frame boundaries or "key frames". By way of example, in MPEG streams, I frames are key frames, P frames are not. See, for further 30 example, Fig. 3, which illustratively depicts a portion 300 of a media stream that includes two key frames 302 and a plurality of other frames 304. As shown, there can be a long rendering time period 306 between key frames 302.

[0058] Consequently, startup logic 210 is advantageously 40 configured to selectively scan through the buffered content to locate, and/or otherwise identify/keep track of, a discrete point at which to start the streaming process for a new client device. Preferably, the discrete point will be the earliest one in buffer 212.

[0059] Propagation latency is another factor for determining where to start sending content from the buffered list. Since essentially old (i.e., buffered) content is sent to new client devices, and the event may be a live broadcast, a time shift is introduced. The magnitude of the 50 time shift resulting from fast startup logic 210 (and media server 208) is related to the amount of buffering done as well as the starting point chosen for content sent to new client devices.

[0060] Clients can randomly connect at any point during 55 a broadcast, and the 10 second buffer list used in this example is constantly changing similar to a "waterfall" or "sliding window". Therefore, the amount of content sent at a rate greater than the encoded bitrate and

the starting point for transmission of content will vary over time. Furthermore, since each client device may have a different link bandwidth, the rate of the accelerated transmission may vary too. Each client device may even have different client-side buffer settings. All of these factors imply that client devices will not be synchronized during the rendering process.

[0061] If a client device connecting to server 202 does not have considerable additional network bandwidth available, sending the earliest usable point in the content buffer list may unnecessarily increase the propagation latency for that specific client device. Thus, server device 202, and more particularly fast startup logic 210, is configured to "balance" the need for minimizing the startup time with the need for minimizing the propagation time. For example, to help balance the conflicting requirements of minimizing propagation latency and startup latency, logic similar to the following can be employed:

```
AccelDuration = RequestedAccelDuration -  
( AccelRate * RequestedAccelDuration);  
RequestedAccelDuration: The requested duration  
of the acceleration.  
AccelRate: The ratio of the (encoded bitrate of the  
content) / (link bandwidth).  
AccelDuration: The amount of content sent from the  
server buffers.
```

[0062] The above exemplary logic essentially reduces the amount of "time-shifted" content sent from server buffer 212 as the encoded bitrate of the content approaches the available link bandwidth.

[0063] For on-demand content, fast startup logic 210 is configured similar to the broadcast scenario except that there is no existing buffer when client device 206 connects. Therefore, fast startup logic 210 builds a buffer list quickly to satisfy the fast startup request. This is possible because a typical media storage device 216 is capable of delivering the on-demand content at a rate that is much faster than what client device 206 is requesting.

[0064] In certain exemplary implementations, the actual protocol mechanism used by client device 206 to request the accelerated buffering involves the use of headers. By way of example, for the RTSP protocol, a header "X-Accelerate-Streaming" is defined, which is used with the PLAY command. This header includes information regarding the client request for the duration of the acceleration and also the bandwidth to use for the acceleration. For example, "AccelDuration=10000;AccelBandwidth=1048576" might be included in a typical "X-Accelerate-Streaming" header by the client. This would inform the server that the client wishes to have 10,000 ms worth of content accelerated at a rate of 1,048,576 bits/s.

[0065] With the HTTP protocol, for example, client fast startup logic 226 can use directives in the commonly

used PRAGMA header in the GET command to specify the fast startup parameters. The text below shows the contents of a sample PRAGMA header in a GET request asking for fast startup.

5 "LinkBW=2147483647, AccelBW=1048576, AccelDuration=10000"

[0066] In this exemplary request, the client fast startup logic 226 is informing server fast startup logic 210 that the link bandwidth is 2,147,483,647 bits/s, but it only 10 wants the content accelerated at a rate of 1,048,576 bits/s for a duration of 10,000 ms.

[0067] Reference is now made to Fig. 4, which is a flow diagram depicting a process 400 suitable for use in client device 206. In step 402, media client 224 connects 15 to media server 208. In step 404, fast startup logic 226 determines the link bandwidth, for example as described above. In step 406, media client 224 asks the media server 210 for information about the available streamable content, including the bandwidth of individual content streams.

[0068] In step 410, fast startup logic 226 requests streamable content from fast startup logic 210. In step 410, fast startup logic 226 selects the initial fast streaming bitrate and the subsequent slower streaming bitrate. 25 Fast startup logic 226 also determines an amount of streamed media that is to be sent at the initial fast streaming bitrate.

[0069] For example, based on buffer 230 settings, the link bandwidth, and the encoded bitrate of the content, 30 fast startup logic 226 can decide whether to submit a request to accelerate the transmission of content in order to fill buffer 230 quickly. If client device 206 decides to request fast startup, custom header syntax can be added to the final command that initiates the delivery of 35 content from server device 202.

[0070] Thereafter, in step 410, media client 224 begins receiving streamed content from media server 208. In step 412, a beginning portion of the content is received at the initial faster streaming bitrate, which is 40 greater than the encoded bitrate. Subsequently, in step 414, further portions of the streamed content are received at the slower streaming bitrate, which is about equal to the encoded bitrate.

[0071] Reference is now made to Fig. 5, which is a 45 flow diagram depicting a process 500 suitable for use in server device 202. In step 502, media server 208 processes a client connect request, and responds to client requests for information about the streamable content. In step 504, fast startup logic 210 responds to a received 50 request from client device 206 for streaming media with fast startup. In step 506, if the requested streamable content is on-demand content, then fast startup logic 210 attempts to satisfy the fast startup request by fetching the necessary content from media storage device 216.

[0072] Alternatively, if the requested streamable content includes broadcast content, then, in step 508, fast startup logic 210 uses the fast startup parameters re-

ceived from fast startup logic 226 to determine at what point in the broadcast content the content can begin streaming from buffer 212. When possible, fast startup logic 210 will preferably start the streaming at discrete starting points in the buffer list so that media client 224 can immediately begin buffering useful content packets.

[0073] Next, in step 510, fast startup logic 210 initially streams the applicable content at the faster streaming bitrate, and subsequently, in step 512, at the lower streaming bitrate.

[0074] Reference is now made to Fig. 6, which is a time line graph 600 illustrating an exemplary transmission bitrate value 602 associated with a requested streaming media program with fast startup. At time t_0 , client device 206 requests the streaming media program. In response server device 202 begins accessing buffered content. At time t_1 , server device 202 begins transmitting the content at a fast streaming bitrate. In this example, at about time t_2 , client device 206 has received and buffered enough streamed media to begin rendering the content. At time t_3 , server device 202 has delivered the requested amount of fast startup data requested by client device 206. As such, the streaming bitrate is reduced to about the encoded bitrate.

[0075] For example, assume that content packets would normally be streamed at a fixed rate of about 56 kbps, even though the link bandwidth for the client device is about 700 kbps. In the fast startup scenario illustrated above, as requested, the content packets that comprise about the first 10 seconds of the media can be streamed at about the link bandwidth rate. Here, this would take roughly 0.8 seconds. Thereafter, the remaining content packets are streamed at the lower encoded bitrate.

[0076] In this example, if the round trip time is reasonably short in duration, then media server 208 will begin the fast startup stream about 0.1 seconds after the request is made. Media client 224 will have received about 5 seconds of the streaming media program at about 0.5 seconds following the initial request, and can begin rendering at about that time. The requested 10 seconds of fast startup streamed media will have been received at about 0.9 seconds following the initial request. At that time, renderer 228 will have rendered about 0.4 seconds of content, and about 9.6 seconds of content will be stored in buffer 230.

[0077] Thus, in this example, the startup time was reduced from over 5 seconds to less than about 1 second. Furthermore, client device 206 will be able to maintain about 10 seconds of buffered content. This additional buffering allows client device 206 to avoid short pauses due for example to jitter and other potentially longer network brownouts, etc.

[0078] Although some preferred implementations of the various methods and apparatuses of the present invention have been illustrated in the accompanying Drawings and described in the foregoing Detailed Description, it will be understood that the invention is not

limited to the exemplary implementations disclosed, but is capable of numerous rearrangements, modifications and substitutions without departing from the spirit of the invention as set forth and defined by the following claims.

Claims

10. 1. A method for use in a client device, the method comprising:
15. identifying streamable media content available from a server device over a communication link, said streamable media content having associated with it an encoded bitrate;
20. providing a startup request to said server device over said communication link, said startup request identifying said streamable media content, a communication link bandwidth, and an amount of said streamable media content to be transmitted at a bitrate greater than said encoded bitrate and no greater than about said communication link bandwidth.
25.
2. The method as recited in Claim 1, further comprising:
30. while receiving said amount of said streamable media content from said server device over said communication link:
35. buffering said received streamable media, and
subsequently rendering said buffered streamed media.
40.
3. The method as recited in Claim 2, wherein receiving said amount of said streamable media content from said server device over said communication link further includes:
45. initially receiving streamed media content associated with a discrete rendering point.
50.
4. The method as recited in Claim 3, wherein said discrete rendering point includes a key frame.
5. The method as recited in Claim 1, wherein said amount of said streamable media content identifies an initial period of playback time.
6. The method as recited in Claim 1, wherein said communication link bandwidth is actively determined by said client device.
55. 7. The method as recited in Claim 1, wherein said streamable media is selected from a group of

streamable media comprising on-demand media and broadcast media.

8. The method as recited in Claim 1, wherein identifying said streamable media content available from said server device includes transmitting at least one session description protocol (SDP) message to said server device.

9. The method as recited in Claim 1, wherein providing said startup request to said server device includes transmitting at least one message formatted according to at least one streaming media protocol selected from a group of protocols comprising MMS, Real Time Streaming Protocol (RTSP), and HyperText Transfer Protocol (HTTP).

10. An apparatus for use in a client device, the apparatus comprising:

memory suitable for buffering streamable media content having associated with it an encoded bitrate; and

logic operatively coupled to said memory and configured to generate at least one startup request that identifies a desired streamable media content, a communication link bandwidth, and an amount of said desired streamable media content that is to be provided at a bitrate greater than said encoded bitrate but no greater than about said communication link bandwidth.

11. The apparatus as recited in Claim 10, further comprising:

a communication interface operatively coupled to said logic and configurable connect to a server device through a communication link, said communication interface being further configurable to transmit said at least one startup message to said server device and receive said amount of said streamable media content from said server device over said communication link; and

wherein said logic is further configured to buffer said received streamable media in said memory, and subsequently render said buffered streamed media.

12. The apparatus as recited in Claim 11, wherein said communication interface is configurable receive initially streamed media content associated with a discrete rendering point.

13. The apparatus as recited in Claim 12, wherein said discrete rendering point includes a key frame.

14. The apparatus as recited in Claim 11, wherein said logic is further configurable to identify said streamable media content available from said server device

5

15. The apparatus as recited in Claim 10, wherein said amount of said streamable media content identifies an initial period of playback time.

10

16. The apparatus as recited in Claim 10, wherein said logic is configured to actively determine said communication link bandwidth.

15

17. The apparatus as recited in Claim 10, wherein said streamable media is selected from a group of streamable media comprising on-demand media and broadcast media.

20

18. A computer-readable medium comprising computer-executable instructions for:

identifying streamable media content available from a server device over a communication link, said streamable media content having associated with it an encoded bitrate;

providing a startup request to said server device over said communication link, said startup request identifying said streamable media content, a communication link bandwidth, and an amount of said streamable media content to be transmitted at a bitrate greater than said encoded bitrate and no greater than about said communication link bandwidth.

25

30

35

19. The computer-readable medium as recited in Claim 18, further comprising computer-executable instructions for:

receiving said amount of said streamable media content from said server device over said communication link;

buffering said received streamable media; and

subsequently rendering said buffered streamed media.

40

45

20. The computer-readable medium as recited in Claim 19, wherein receiving said amount of said streamable media content from said server device over said communication link further includes:

initially receiving streamed media content associated with a discrete rendering point.

50

55

21. The computer-readable medium as recited in Claim 20, wherein said discrete rendering point includes a key frame.

22. The computer-readable medium as recited in Claim

18, wherein said amount of said streamable media content identifies an initial period of playback time.

23. The computer-readable medium as recited in Claim 18, further comprising computer-executable instructions for:

actively determining said communication link bandwidth.

24. A method for use in a server device, the method comprising:

identifying streamable media content available to a client device over a communication link, said streamable media content having associated with it an encoded bitrate; receiving a startup request from said client device over said communication link, said startup request specifying said streamable media content to be transmitted to said client, a communication link bandwidth, and an amount of said streamable media content to be transmitted at a bitrate greater than said encoded bitrate and no greater than about said communication link bandwidth.

25. The method as recited in Claim 24, further comprising:

buffering at least said amount of said streamable media content; and transmitting said amount of said buffered streamable media content to said client device at said bitrate greater than said encoded bitrate and no greater than about said communication link bandwidth.

26. The method as recited in Claim 25, further comprising:

after transmitting said amount of said buffered streamable media content, transmitting subsequent portions of said streamable media content to said client device at a bitrate about equal to said encoded bitrate.

27. The method as recited in Claim 25, wherein transmitting said amount of said buffered streamable media content to said client device further includes:

locating a discrete rendering point in said amount of said buffered streamable media content; and initiating transmission of said amount of said buffered streamable media content beginning with said discrete rendering point.

28. The method as recited in Claim 27, wherein said discrete rendering point includes a key frame.

29. The method as recited in Claim 25, wherein said amount of said streamable media content identifies an initial period of playback time.

30. The method as recited in Claim 25, wherein said streamable media is selected from a group of streamable media comprising on-demand media and broadcast media.

31. An apparatus for use in a server device, the apparatus comprising:

memory suitable for buffering streamable media content having associated with it an encoded bitrate; and logic operatively coupled to said memory and configurable to respond in accord with a startup request, said startup request specifying streamable media content to be provided to said client, a communication link bandwidth associated with a communications link established between said server device and said logic, and an amount of said streamable media content to be transmitted at a bitrate greater than said encoded bitrate and no greater than about said communication link bandwidth.

32. The apparatus as recited in Claim 31, wherein said logic is configured to buffer at least said amount of said streamable media content in said memory, and transmit said amount of said buffered streamable media content to said client device at said bitrate greater than said encoded bitrate but no greater than about said communication link bandwidth.

33. The apparatus as recited in Claim 32, wherein said logic is further configured to, after transmitting said amount of said buffered streamable media content, transmit subsequent portions of said streamable media content to said client device at a bitrate about equal to said encoded bitrate.

34. The apparatus as recited in Claim 32, wherein, in transmitting said amount of said buffered streamable media content to said client device, said logic locates a discrete rendering point in said amount of said buffered streamable media content and initiates transmission of said amount of said buffered streamable media content beginning with said discrete rendering point.

35. The apparatus as recited in Claim 34, wherein said discrete rendering point includes a key frame.

36. The apparatus as recited in Claim 31, wherein said

amount of said streamable media content identifies an initial period of playback time.

37. The apparatus as recited in Claim 31, wherein said streamable media is selected from a group of streamable media comprising on-demand media and broadcast media.

38. A computer-readable medium comprising computer-executable instructions for:

identifying streamable media content available to a client device over a communication link, said streamable media content having associated with it an encoded bitrate; receiving a startup request from said client device over said communication link, said startup request specifying said streamable media content to be transmitted to said client, a communication link bandwidth, and an amount of said streamable media content to be transmitted at a bitrate greater than said encoded bitrate and no greater than about said communication link bandwidth.

39. The computer-readable medium as recited in Claim 38, further comprising computer-executable instructions for:

buffering at least said amount of said streamable media content; and transmitting said amount of said buffered streamable media content to said client device at said bitrate greater than said encoded bitrate and no greater than about said communication link bandwidth.

40. The computer-readable medium as recited in Claim 39, further comprising computer-executable instructions for:

after transmitting said amount of said buffered streamable media content, transmitting subsequent portions of said streamable media content to said client device at a bitrate about equal to said encoded bitrate.

41. The computer-readable medium as recited in Claim 39, wherein transmitting said amount of said buffered streamable media content to said client device further includes:

locating a discrete rendering point in said amount of said buffered streamable media content; and initiating transmission of said amount of said buffered streamable media content beginning with said discrete rendering point.

42. The computer-readable medium as recited in Claim 41, wherein said discrete rendering point includes a key frame.

5 43. The computer-readable medium as recited in Claim 38, wherein said amount of said streamable media content identifies an initial period of playback time.

10 44. The computer-readable medium as recited in Claim 38, wherein said streamable media is selected from a group of streamable media comprising on-demand media and broadcast media.

15 45. A system comprising:

a network configured to provide a communication link; a server device operatively coupled to said network; and a client device operatively coupled to said network, and configured to send at least one startup request to said server device over said communication link, said startup request identifying a streamable media content to be provided to said client device over said communication link, a client determined communication link bandwidth associated with said communication link, and an amount of said desired streamable media content that to be provided at a bitrate greater than said encoded bitrate but no greater than about said communication link bandwidth.

20 46. The system as recited in Claim 45, wherein said server device is configured to buffer at least said amount of said streamable media content, and transmit said amount of said buffered streamable media content to said client device at said bitrate greater than said encoded bitrate but no greater than about said communication link bandwidth.

25 47. The system as recited in Claim 46, wherein said server device is further configured to, after transmitting said amount of said buffered streamable media content, transmit subsequent portions of said streamable media content to said client device at a bitrate about equal to said encoded bitrate.

30 48. The system as recited in Claim 46, wherein, in transmitting said amount of said buffered streamable media content to said client device said server device locates a discrete rendering point in said amount of said buffered streamable media content and initiates transmission of said amount of said buffered streamable media content beginning with said discrete rendering point.

35 49. The system as recited in Claim 48, wherein said discrete rendering point includes a key frame.

40

45

50

55

50. The system as recited in Claim 45, wherein said amount of said streamable media content identifies an initial period of playback time.

51. The system as recited in Claim 45, wherein said streamable media is selected from a group of streamable media comprising on-demand media and broadcast media.

52. The system as recited in Claim 45, wherein said client device is further configured to buffer received streamable media content, and subsequently render said buffered streamed media content. 10

15

20

25

30

35

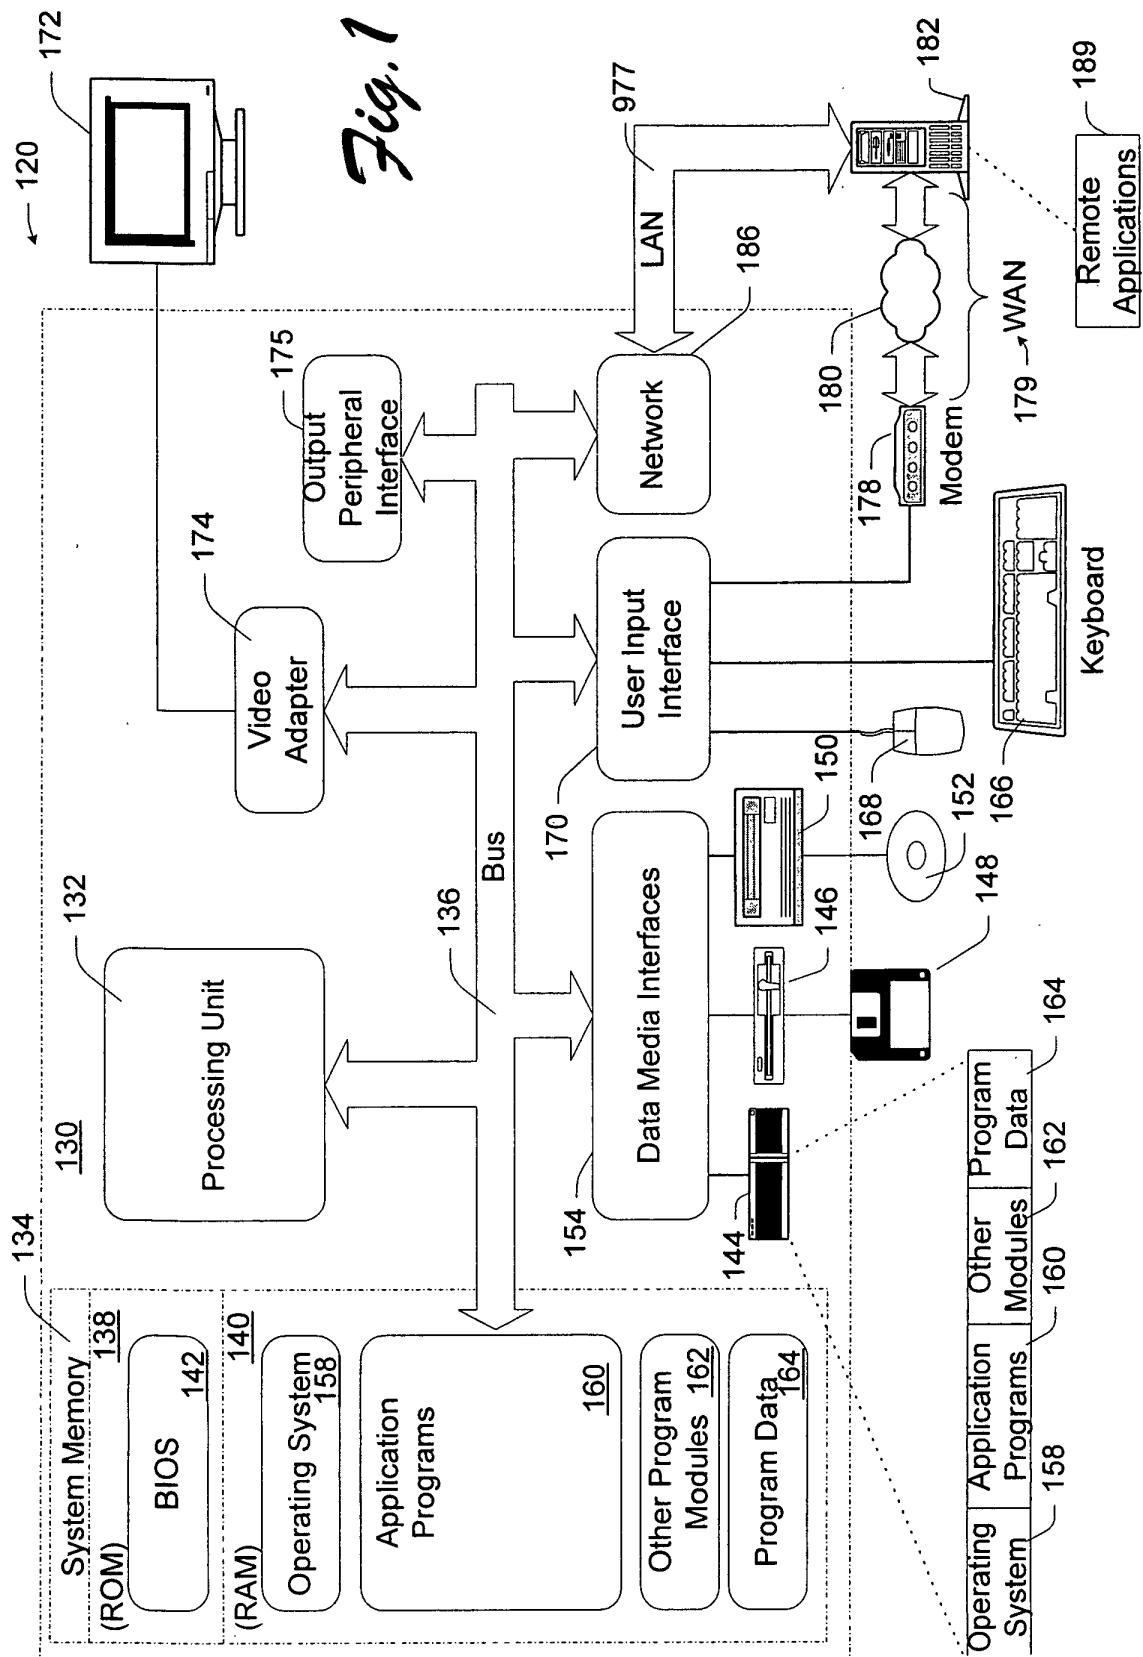
40

45

50

55

Zig. 1



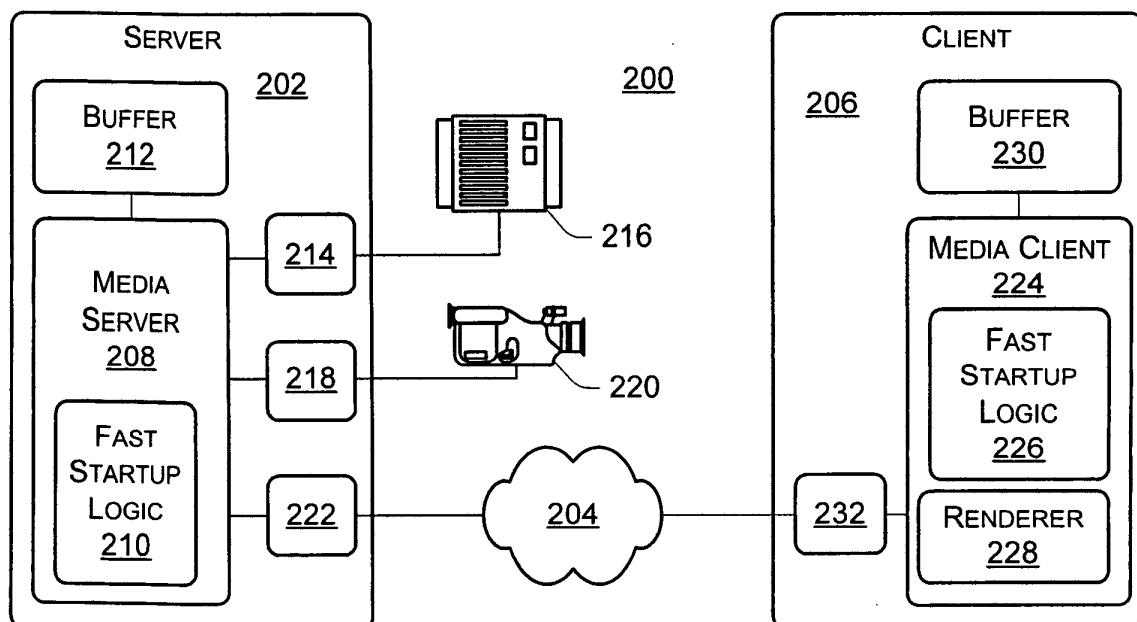


Fig. 2

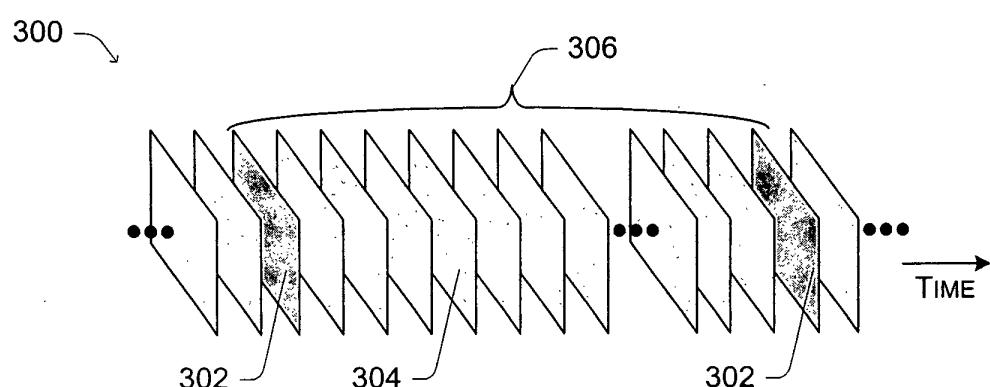


Fig. 3

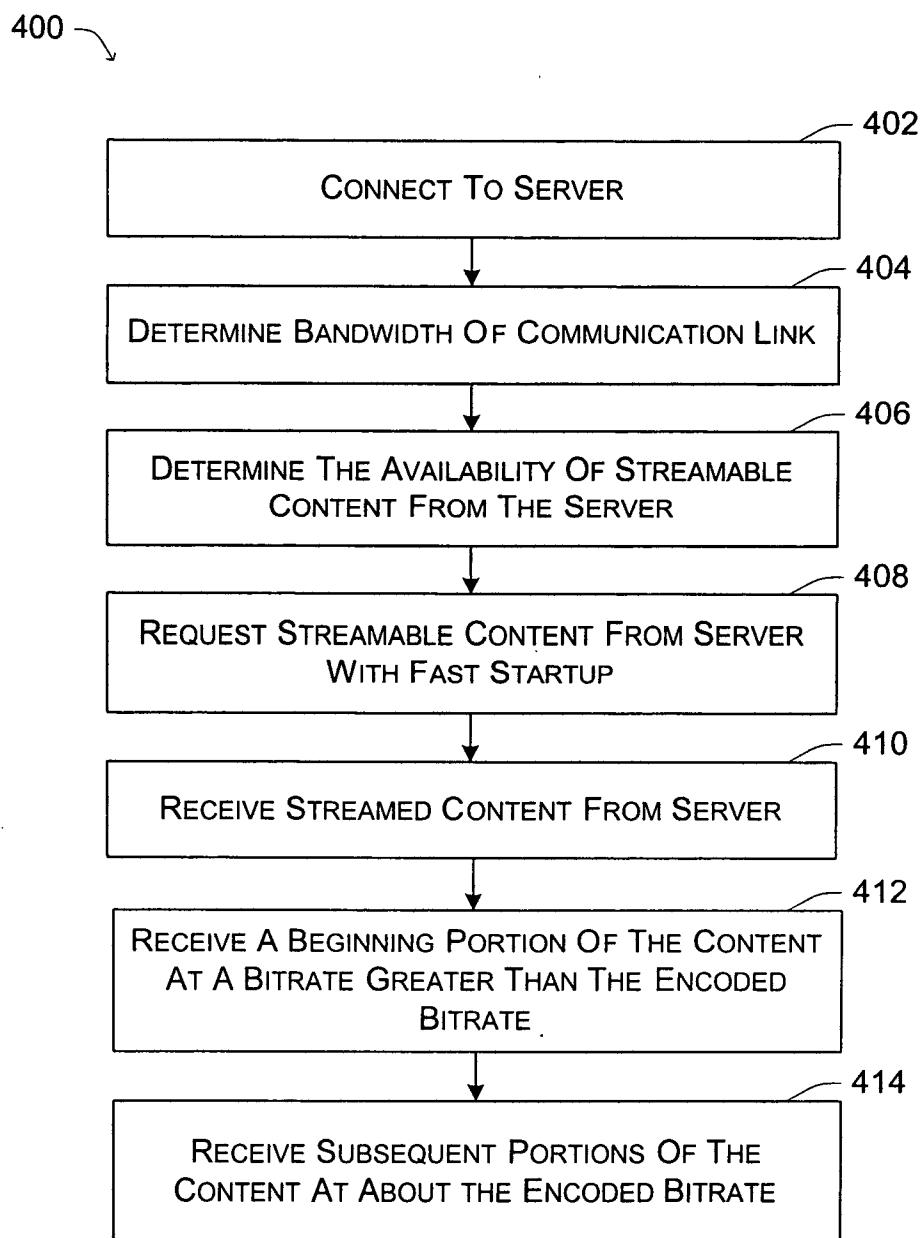


Fig. 4

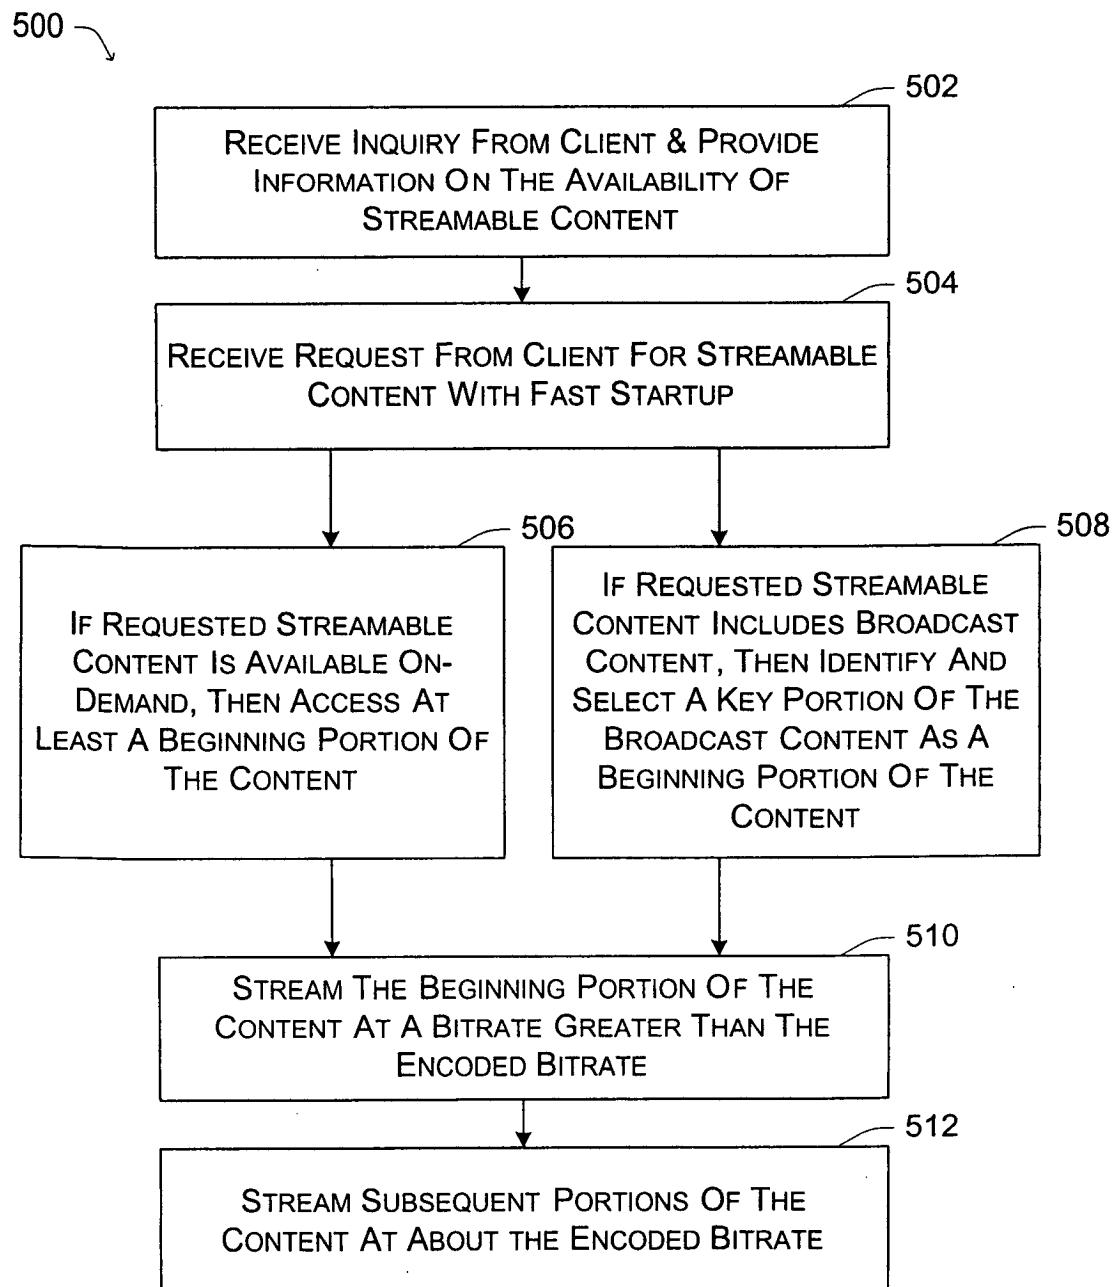


Fig. 5

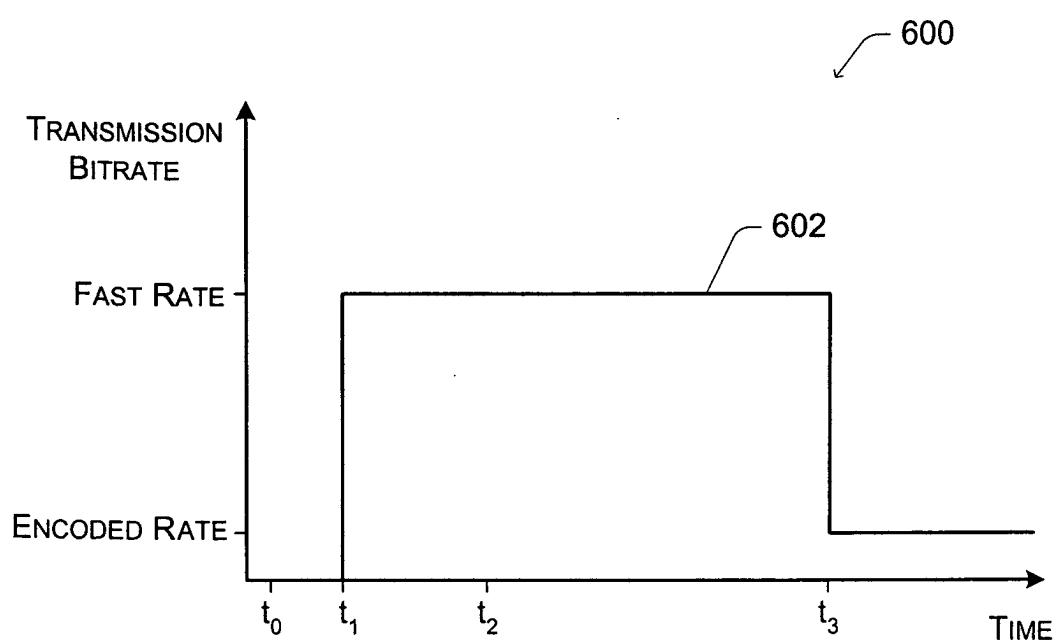


Fig. 6